The AMD processor roadmap for industry-standard servers

technology brief, 3rd Edition

Abstract

In 2003, AMD introduced the AMD Opteron™ processor with 64-bit extensions. Opteron was the first x86-based processor capable of supporting 64-bit operating systems and applications. AMD introduced dual-core Opteron processors in 2005.

This paper discusses current and near future Opteron processors, the high-level architecture of the Opteron family, and how the processors are used in ProLiant server products.

For additional details about Opteron architecture, see the technology brief titled "Characterizing x86 processors for industry-standard servers: AMD Opteron and Intel Xeon." [1]

Introduction

The Opteron™ family of processors is AMD's offering for the industry-standard server market.

HP is providing AMD processors in the ProLiant server product line to offer enterprise customers expanded options for improved performance while maintaining cost-effective infrastructures.

AMD Opteron processors feature Direct Connect Architecture, AMD64, AMD Virtualization, and AMD PowerNow!. In some ProLiant models, the installed AMD Opteron processors can be upgraded to Quad-Core processors when they become available.

Key Opteron features

All Opteron processors to date use AMD K8 architecture and share three key features:

- AMD64 technology
- Integrated memory controller
- Direct Connect I/O Architecture

AMD64 technology

Introduced in 2003, AMD64 is the AMD micro-architecture and instruction set that provides full support for 64-bit operating systems and applications.

AMD64 is a set of logical extensions to the familiar 32-bit x86 instruction set. It includes the basic x86, SSE (Single SIMD Extensions), SSE2, and SSE3 instruction sets. It is fully backward compatible with legacy 32-bit software, and it can run both 32-bit and 64-bit applications simultaneously in separate processes, without emulation.

The most important feature of AMD64 is support for very large virtual and physical memory in a flat address space. In fact, the full 64-bit space is much larger than current systems require. As a result, current Opteron implementations use 40 bits for physical memory (1 terabyte) and 48 bits for virtual memory (256 Terabytes).

All the registers within Opteron are at least 64 bits wide and all instructions have a 64-bit operand option, making Opteron a true 64-bit processor with enough addressing capability to perform the largest computing tasks in the industry-standard server market.

1 http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00238028/c00238028.pdf

Integrated memory controller

The Opteron processor contains an on-die integrated dual-channel SDRAM memory controller that operates at the core speed of the processor. (Memory controllers for the recently introduced Revision F processor support DDR2, while previous revisions support DDR.)

Integrating the controller into the processor means that memory performance can scale linearly based on the number of processors in a multi-processor system. For example, in a multi-processor system, the integrated memory controller allows for multiple memory requests in parallel, thereby increasing the effective memory bandwidth and decreasing average memory latency.

The memory controller has a 128-bit interface that is capable of supporting up to eight DDR2 DIMMs (with DDR2-667, four DIMMs per channel with memory bandwidth up to 10.7 GB/s per processor).

Direct Connect I/O Architecture

The Direct Connect I/O Architecture uses the HyperTransport™ bus to directly connect the CPU to DRAM main memory and other CPUs.

HyperTransport is a point-to-point interconnect with dual uni-directional links (see Figure 1). Compared to a shared (bi-directional) parallel front-side bus, a point-to-point interconnect has the advantage of no overhead for bus arbitration and easier signal integrity maintenance, resulting in a scalable, high-bandwidth architecture. HyperTransport uses double data rate to transfer 16 bits of data on both the rising and falling edge of the clock signals, resulting in an effective 32 bits of data per clock cycle on the 16-bit link.

Opteron processors use a HyperTransport clock speed of 1 GHz. This provides an effective operating frequency of up to 2000 MT/s (megatransfers per second) and an effective transfer rate of up to 4 GB/s in each direction. Since transfers can occur in both directions simultaneously, an aggregate transfer rate of 8 GB/s can be achieved in 16-bit HyperTransport I/O links.

Figure 1. HyperTransport bus with dual uni-directional links, one upstream and one downstream.

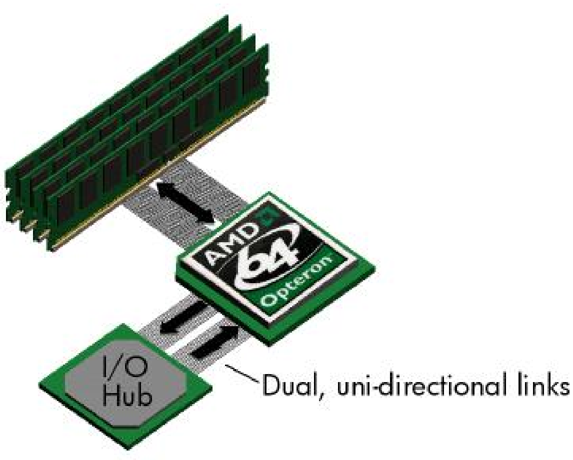

Multi-Core technology

To improve processor performance, the most common solutions have been to increase core frequency and/or cache size. However, both of these solutions increase power consumption (and heat generation) and have other limitations. Alternatively, higher performance can be achieved by using multiple execution cores per processor. Multi-core processors run applications more efficiently and allow multi-threaded software to achieve higher performance, while maintaining a similar power budget to single-core processors. Today, multi-core technology includes dual-core and quad-core processors. Even more cores per processor are likely in the future.

AMD introduced its first dual-core AMD64 processor in 2005; it was manufactured using a 90 nm process. The Opteron processor is essentially divided into two parts: execution and communications, with a system request interface and crossbar switch linking these two parts (see Figure 2). The crossbar switch architecture enabled Opteron to transition easily from single-core to dual-core processors without fundamental design changes.

Each execution core includes a 64-KB L1 data cache, a 64-KB L1 instruction cache, and a 1-MB L2 cache. The system request interface manages and prioritizes the processor requests to the crossbar switch. The crossbar switch connects both processor cores directly to communications: I/O (via HyperTransport links) and the memory controller. The memory controller and the HyperTransport links remain the same as in a single core system.

Figure. 2. Socket designs—Socket F (1207) and Socket AM2—to support dual-core AMD Revision F processors.

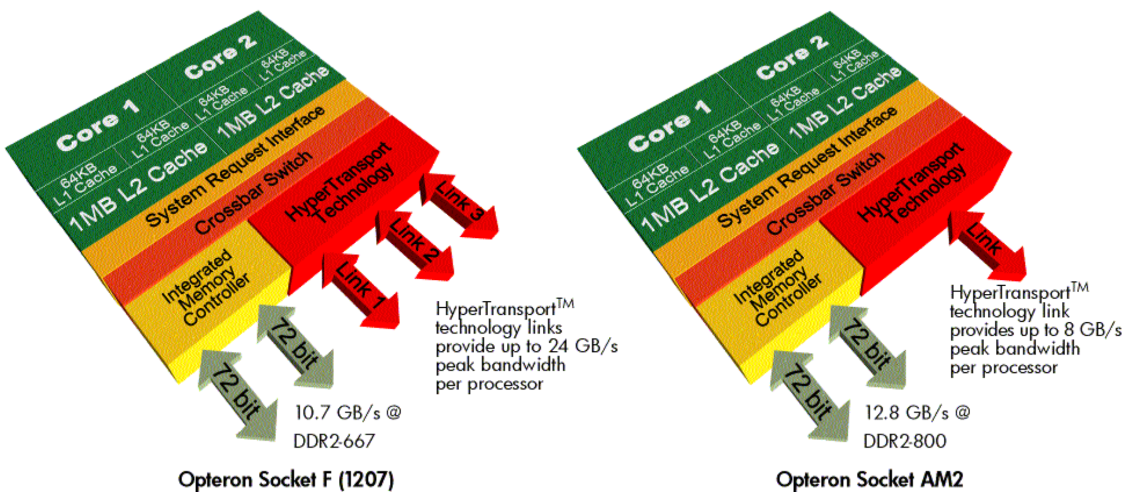

The primary difference between the processors designed for single, dual, or multi-core systems is in the way the processor uses the HyperTransport link(s). As described in the "Naming conventions" section of this paper, there are 1000, 2000, and 8000 series Opteron processors. In the 1000 series, the single HyperTransport link can only connect to I/O in a non-coherent link. This means that the 1000 series Opteron processors are limited to single-processor systems. In the 2000 series, one of three HyperTransport links can connect to one other Opteron processor in a coherent link. The other links can connect to I/O (non-coherent link); thus, 2000 series Opteron processors can

be used in dual-processor systems. With the 8000 series Opteron processors, all three HyperTransport links can connect to other Opteron processors or to I/O.

Multi-core solutions become increasingly attractive with reduction in features size (for example, from 90 nm to 65 nm to 45 nm). Smaller cores require less power, permitting more cores to be built into a single processor. AMD's 65-nm process will allow quad-core processors (Barcelona) to be built. Quad-core Opteron processors (first demonstrated in November 2006) are predicted to reach the market in the second half of 2007. AMD indicates that this quad-core processor will provide 40 to 70 percent greater performance than its current dual-core processors. For more information about multi-core processors, see the AMD whitepaper titled "Multi-Core Processors—The Next Evolution in Computing."[2]

Current Opteron Revisions

Opteron processor "generations" are called Revisions. The current generation, Revision F (Rev F), was introduced in 2006. A number of features from previous revisions remain unchanged in the Rev F processor:

- 64-KB/64-KB data/instruction L1 cache per core
- 1-MB L2 cache per core
- 1-GHz HyperTransport

Key improvements in Revision F include:

- DDR2 memory support
- Hardware assisted virtualization (AMD V™)
- Power management (PowerNow!) improvements
- Quad-core upgradeability

DDR2 memory requires up to 30 percent less power than DDR.

AMD Virtualization™ hardware assistance directly supports virtualization with industry-standard servers, which reduces complexity and improves performance.

PowerNow! Technology with Optimized Power Management reduces power requirements and heat generation by reducing the processor's clock speed and voltage during periods when the CPU is not fully utilized. Up to five power states are supported. Power consumption at idle is reduced by up to 75 percent.

Quad-core upgradeability means that Revision F sockets are pin-compatible with and will support quad-core processors within the same power and thermal envelopes.

[2] See the AMD whitepaper "Multi-Core Processors—The Next Evolution in Computing" at, http://multicore.amd.com/Resources/33211A_Multi-Core_WP_en.pdf

Naming conventions

Current Revision F (DDR2) Opteron processors have a four-digit model number, XYZZ, where X designates scalability, Y socket generation, and ZZ relative performance[3]. The first digit, X, indicates the processor series, where 1 indicates 1000 series, 2 indicates 2000 series, and 8 indicates 8000 series:

- 1000 series – For single processor servers; no HyperTransport links can be used to connect to other processors
- 2000 series – For dual processor servers; only one HyperTransport link can be used to connect to other processors
- 8000 series – For up to eight-processor servers; all three HyperTransport links can be used to connect to other processors

The second digit, Y, indicates socket generation, where 2 indicates Socket AM2 or Socket F (1207). Socket AM2 supports Series 1000 processors; Socket F (1207) supports Series 2000 and 8000 processors.

The last two digits, ZZ, indicate the relative performance within the series. Higher numbers indicate higher performance.

In addition, the model number can include a suffix designator to indicate a non-standard power level. HE designates a lower power version, and SE a higher power version. For example, Model 2220, Model 2220 HE, and Model 2220 SE all offer equivalent performance, but differ in power consumption.

Older processors used a three-digit model number, XZZ, where X designates scalability and ZZ relative performance.

- 100 series—For single-processor servers
- 200 series—For dual-processor servers
- 800 series—For up to eight processor servers

Socket 939 supports Series 100 processors, and Socket 940 supports Series 200 and 800 processors. The AMD website includes a quick reference guide[4] that details each processor part number by socket, revision (stepping), core frequency, manufacturing process (90 nm or 130 nm), HyperTransport frequency, and wattage.

Software licensing

Customers should be aware of possible changes in software licensing for use of multi-core processors. At this writing, major OS vendors, such as Microsoft, currently treat multi-core processors as performance improvements to a single processor[5]; they are not making a distinction for licensing purposes among processors with one, two, four, or more cores. However, customers should check with their OS and application vendors to determine their particular licensing requirements.

[3] http://www.amd.com/us-en/Processors/ProductInformation/0,,30_118_8796_14266,00.html
[4] http://www.amdcompare.com/us-en/opteron/
[5] Refer to the Microsoft website http://www.microsoft.com/licensing/highlights/multicore.mspx

Future improvements

The Revision H Opteron is scheduled for introduction in the second half of 2007. It will include:

- New core micro-architecture – K8L
- True quad-core on a single die
- Extensions to AMD64 instruction set – SSE and bit manipulation
- 128-bit FPU for improved floating point and graphics performance
- Increased instructions-per-cycle (IPC)
- DDR2 memory
- Dedicated 64-KB L1 cache and 512-KB L2 cache for each core
- New 2-MB or larger L3 cache shared among all cores
- Independent power control for each core depending on workload
- Significant performance per watt improvements

AMD plans to migrate Opteron production from its current 90 nm process to a 65 nm process in 2007 and to a 45 nm process in 2008. Revision H is expected to be introduced with a 65 nm process (see Figure 3).

Figure 3. Quad-Core Opteron™ processor design for Socket F (1207)

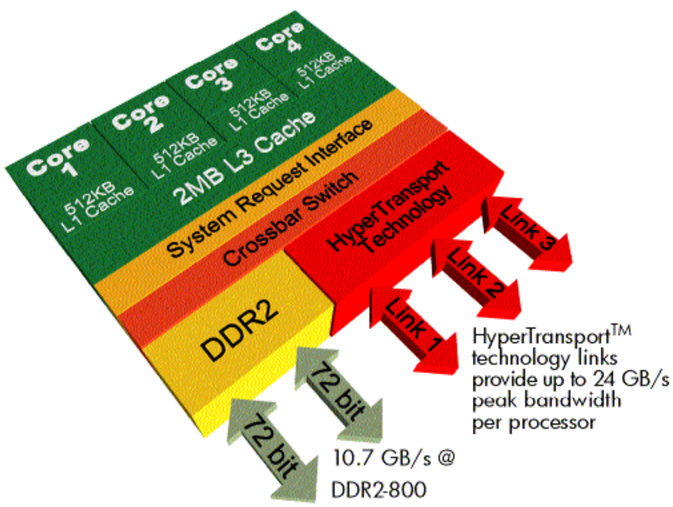

AMD has announced the Torrenza initiative, which will provide an additional chip socket for a co-processor on the motherboard. This socket will include a HyperTransport bus connection, and will support graphics and other more specialized third-party co-processors.

AMD has announced plans to support DDR3 in 2008, which will provide higher performance and lower power requirements.

Conclusion

HP ProLiant servers using the Opteron processor family have proven their performance in numerous benchmarks and systems. They provide a straightforward progression from 32-bit to 64-bit computing.

HP ProLiant servers continue to offer both AMD Opteron and Intel® Xeon™ processor architectures to deliver the best possible choice to customers. The addition of multi-core AMD Opteron technology will take advantage of multi-threaded applications and reduce latencies, while allowing the processors to stay within the same power budgets as single-core versions.

AMD is expected to continue improving the Opteron processor family with faster memory and HyperTransport speeds.

For more information

For additional information, refer to the resources listed below.

HyperTransport Consortium	http://www.hypertransport.org/index.cfm
Multi-Core Processors—The Next Evolution in Computing	http://multicore.amd.com/Resources/33211A_Multi-Core_WP_en.pdf
Power Regulator for ProLiant Servers	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00593374/c00593374.pdf
ISS Technology Papers	http://www.hp.com/servers/technology

Call to action

Send comments about this paper to TechCom@HP.com.

TC070703TB, July 2007